# EAST Search History

| Ref # | Hits | Search Query | DBs | Default Operator | Plurals | Time Stamp |
|---|---|---|---|---|---|---|
| S1 | 59 | matthew near huras.in. | US-PGPUB; USPAT | OR | ON | 2007/06/11 11:14 |
| S2 | 28 | sam near lightstone.in. | US-PGPUB; USPAT | OR | ON | 2007/06/11 11:23 |
| S3 | 9 | sujay near parekh.in. | US-PGPUB; USPAT | OR | ON | 2007/06/11 11:24 |
| S4 | 12 | kevin near rose.in. | US-PGPUB; USPAT | OR | ON | 2007/06/11 13:49 |
| S5 | 70005 | (ibm or "international business machines").as. | US-PGPUB; USPAT | OR | ON | 2007/06/11 11:25 |
| S7 | 68 | S5 and (throttl$4).clm. | US-PGPUB; USPAT | OR | ON | 2007/06/11 11:26 |
| S8 | 3163 | 717/124-135.ccls. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 11:34 |
| S9 | 29 | S8 and throttl$4 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 11:34 |
| S10 | 20 | S9 and (@pd<"20031019" or @ad<"20031019" or @prad<"20031019" or @rlad<"20031019") | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/12 08:37 |
| S11 | 207 | S8 and ((limit$4 or reduc$4) near3 resource) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 13:08 |
| S12 | 142 | S11 and performance | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 13:11 |

| | | | | | | |
|---|---|---|---|---|---|---|
| S13 | 113 | S12 and (@pd<"20031019" or @ad<"20031019" or @prad<"20031019" or @rlad<"20031019") | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 15:12 |
| S14 | 341 | throttl$4 adj level | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 13:36 |
| S15 | 14 | S14 and (manag$5 near3 performance) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 13:40 |
| S16 | 4 | ("4807161" \| "5428789" \| "6205468" \| "6834386").PN. | US-PGPUB; USPAT; USOCR | OR | ON | 2007/06/11 13:43 |
| S17 | 100 | john near douceur.in. | US-PGPUB; USPAT | OR | ON | 2007/06/11 13:49 |
| S18 | 3 | S17 and throttl$4 | US-PGPUB; USPAT | OR | ON | 2007/06/11 13:50 |
| S19 | 50 | ("7137019" "20040221184" "20050086659" "20070074057" "5842024" "6032172" "6287765" "6292830" "20040172408" "5435138" "5517977" "5524596" "5754852" "4417131" "5008560" "5212915" "5247682" "5407152" "5485628" "5511872" "5535415" "5688168" "5740650" "5745751" "5746034" "5746035" "5779609" "5784843" "5809708" "5886295" "5899025" "5899036" RE36226 "6067762" "6073399" "6128873" "6202374" "6279643" "6313549" "6356822" "6415567" "6427101" "6430486" "6438545" "6564201" "6748707" "6931795" "20040236620" "20050038721" "5375242" ).pn. | US-PGPUB; USPAT | OR | ON | 2007/06/11 14:02 |
| S20 | 3036 | 718/102-104.ccls. | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 15:12 |

| S21 | 74 | S20 and throttl$4 | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 15:12 |
|-----|------|-------------------|------|------|------|------------------|
| S22 | 68 | S21 and (@pd<"20031019" or @ad<"20031019" or @prad<"20031019" or @rlad<"20031019") | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/11 15:12 |
| S23 | 10 | (US-20070074057-$ or US-20040221184-$ or US-20030088605-$ or US-20050071834-$).did. or (US-7137019-$ or US-7093253-$ or US-5903759-$ or US-6834386-$ or US-7086059-$ or US-6766349-$).did. | US-PGPUB; USPAT | OR | ON | 2007/06/12 08:28 |
| S24 | 0 | S23 and ((reduc$4 or lower$4) near3 (memory or space)) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/12 08:29 |
| S25 | 2337 | performance and throttl$4 and ((reduc$4 or lower$4) near3 (memory or space)) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/12 08:35 |
| S26 | 695 | performance and throttl$4 and ((reduc$4 or lower$4) near3 memory) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/12 08:35 |
| S29 | 31 | S26 and ("717".clas. or "718".clas.) | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/12 08:36 |
| S30 | 27 | S29 and (@pd<"20031019" or @ad<"20031019" or @prad<"20031019" or @rlad<"20031019") | US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/06/12 08:37 |

**PØRTAL**
USPTO

Search:   ⊙ The ACM Digital Library   ○ The Guide

"managing performance" throttling                              **SEARCH**

THE ACM DIGITAL LIBRARY

Feedback  Report a problem  Satisfaction survey

Terms used **managing performance throttling**                **Found 420 of 203,282**

Sort results by   [relevance ▾]        ❤ Save results to a Binder        Try an Advanced Search
                                        ⍰ Search Tips                    Try this search in The ACM Guide
Display results   [expanded form ▾]    ☐ Open results in a new window

Results 1 - 20 of 200          Result page: **1**  2  3  4  5  6  7  8  9  10   next
Best 200 shown                                                 Relevance scale ☐ ▭ ◩ ◪ ■

**1**  **Mitigating Amdahl's Law through EPI Throttling**                              ▭
Murali Annavaram, Ed Grochowski, John Shen
May 2005  **ACM SIGARCH Computer Architecture News , Proceedings of the 32nd**
          **Annual International Symposium on Computer Architecture ISCA '05,**
          Volume 33 Issue 2
**Publisher:** IEEE Computer Society, ACM Press
Full text available: 📄 pdf(202.69 KB)   Additional Information: full citation, abstract, citings, index terms

> This paper is motivated by three recent trends in computer design. First, chip multi-
> processors (CMPs) with increasing numbers of CPU cores per chip are becoming common.
> Second, multi-threaded software that can take advantage of CMPs will soon become
> prevalent. Due to the nature of the algorithms, these multi-threaded programs inherently
> will have phases of sequential execution; Amdahlÿs law dictates that the speedup of such
> parallel programs will be limited by the sequential portion of the comp ...

**2**  **Defending against distributed denial-of-service attacks with max-min fair server-**   ▭
**centric router throttles**
David K. Y. Yau, John C. S. Lui, Feng Liang, Yeung Yam
February 2005  **IEEE/ACM Transactions on Networking (TON),** Volume 13 Issue 1
**Publisher:** IEEE Press

Full text available: 📄 pdf(820.35 KB)   Additional Information: full citation, abstract, references, citings, index terms, review

> Our work targets a network architecture and accompanying algorithms for countering
> distributed denial-of-service (DDoS) attacks directed at an Internet server. The basic
> mechanism is for a server under stress to install a router throttle at selected upstream
> routers. The throttle can be the leaky-bucket rate at which a router can forward packets
> destined for the server. Hence, before aggressive packets can converge to overwhelm the
> server, participating routers proactively regulate the contribut ...

> **Keywords**: congestion control, distributed denial of service, network security, router
> throttling

**3**  **Speculative parallel simulation with an adaptive throttle scheme**                ▭
Seng Chuan Tay, Yong Meng Teo, Siew Theng Kong
June 1997  **ACM SIGSIM Simulation Digest , Proceedings of the eleventh workshop on**
          **Parallel and distributed simulation PADS '97,** Volume 27 Issue 1
**Publisher:** IEEE Computer Society, ACM Press
Full text available: 📄 pdf(939.88 KB)   Additional Information: full citation, abstract, references, citings, index terms
📄 Publisher Site

Excessive rollback recoveries due to overoptimistic event execution in Time Warp simulators often degrade their runtime performance. This paper presents a two-sided throttling scheme to dynamically adjust the event execution speed of Time Warp simulators. The proposed throttle is based on a new concept called global progress window, which allows the individual simulation process to be positioned on a global time scale, thereby to accelerate or suspend their event execution. As each simulation pr ...

4 **Burst reduction properties of rate-control throttles: downstream queue behavior**

Zhen Liu, Don Towsley

February 1995 **IEEE/ACM Transactions on Networking (TON)**, Volume 3 Issue 1

**Publisher:** IEEE Press

Full text available: pdf(1.34 MB)     Additional Information: full citation, references, index terms

5 **Estimating the cost of throttled execution in time warp**

Samir R. Das

July 1996 **ACM SIGSIM Simulation Digest , Proceedings of the tenth workshop on Parallel and distributed simulation PADS '96**, Volume 26 Issue 1

**Publisher:** IEEE Computer Society, ACM Press

Full text available: pdf(442.79 KB)     Additional Information: full citation, abstract, references, citings, index terms

Publisher Site

Over-optimistic execution has long been identified as a major performance bottleneck in Time Warp based parallel simulation systems. An appropriate throttle or control of optimism can improve performance by reducing the number of rollbacks. However, the design of an appropriate throttle is a difficult task, as correct computations on the critical path may be blocked, thus increasing the overall execution time. In this paper we build a cost model for throttled execution that involves both rollbac ...

**Keywords**: discrete event simulation, parallel simulation, time warp, performance model

6 **Efficient network and I/O throttling for fine-grain cycle stealing**

Kyung D. Ryu, Jeffrey K. Hollingsworth, Peter J. Keleher

November 2001 **Proceedings of the 2001 ACM/IEEE conference on Supercomputing (CDROM) Supercomputing '01**

**Publisher:** ACM Press

Full text available: pdf(127.89 KB)     Additional Information: full citation, abstract, references, citings, index terms

This paper proposes and evaluates a new mechanism, rate windows, for I/O and network rate policing. The goal of the proposed system is to provide a simple, yet effective way to enforce resource limits on target classes of jobs in a system. This work was motivated by our Linger Longer infrastructure, which harvests idle cycles in networks of workstations. Network and I/O throttling is crucial because Linger Longer can leave guest jobs on non-idle nodes and machine owners should not be adversely a ...

7 **Instruction flow-based front-end throttling for power-aware high-performance processors**

Amirali Baniasadi, Andreas Moshovos

August 2001 **Proceedings of the 2001 international symposium on Low power electronics and design ISLPED '01**

**Publisher:** ACM Press

Full text available: pdf(141.43 KB)     Additional Information: full citation, references, citings, index terms

8 **Power awareness: Combining compiler and runtime IPC predictions to reduce energy in next generation architectures**

Saurabh Chheda, Osman Unsal, Israel Koren, C. Mani Krishna, Csaba Andras Moritz
April 2004 **Proceedings of the 1st conference on Computing frontiers CF '04**
**Publisher:** ACM Press
Full text available: pdf(336.02 KB)    Additional Information: full citation, abstract, references, index terms

Next generation architectures will require innovative solutions to reduce energy
consumption. One of the trends we expect is more extensive utilization of compiler
information directly targeting energy optimizations. As we show in this paper, static
information provides some unique benefits, not available with runtime hardware-based
techniques alone. To achieve energy reduction, we use IPC information at various
granularities, to adaptively adjust voltage and speed, and to throttle the fetch rat ...

**Keywords**: adaptive voltage scaling, compiler architecture interaction, fetch throttling,
instruction level parallelism, low power design·

**9** Thermal Modeling, Characterization and Management of On-Chip Networks
Li Shang, Li-Shiuan Peh, Amit Kumar, Niraj K. Jha
December 2004 **Proceedings of the 37th annual IEEE/ACM International Symposium
on Microarchitecture MICRO 37**
**Publisher:** IEEE Computer Society
Full text available: pdf(551.38 KB)    Additional Information: full citation, abstract, citings

Due to the wire delay constraints in deep submicron technology and increasing demand
for on-chip bandwidth, networks are becoming the pervasive interconnect fabric to
connect processing elements on chip. With ever-increasing power density and cooling
costs, the thermal impact of on-chip networks needs to be urgently addressed. In this
work, we first characterize the thermal profile of the MIT Raw chip. Our study shows
networks having comparable thermal impact as the processing elements and contr ...

**10** Performance bonds for flow control protocols
Rajeev Agrawal, Rene L. Cruz, Clayton Okino, Rajendran Rajan
June 1999 **IEEE/ACM Transactions on Networking (TON)**, Volume 7 Issue 3
**Publisher:** IEEE Press
Full text available: pdf(298.23 KB)    Additional Information: full citation, references, citings, index terms

**Keywords**: adaptive service, burstiness, delay, guaranteed service, network calculus,
queueing, regulator, scheduler, service curve

**11** Chiron-1: a software architecture for user interface development, maintenance, and
run-time support
Richard N. Taylor, Kari A. Nies, Gregory Alan Bolcer, Craig A. MacFarlane, Kenneth M.
Anderson, Gregory F. Johnson
June 1995 **ACM Transactions on Computer-Human Interaction (TOCHI)**, Volume 2 Issue 2
**Publisher:** ACM Press
Full text available: pdf(2.65 MB)    Additional Information: full citation, abstract, references, citings, index
terms, review

The Chiron-1 user interface system demonstrates key techniques that enable a strict
separation of an application from its user interface. These techniques include separating
the control-flow aspects of the application and user interface: they are concurrent and
may contain many threads. Chiron also separates windowing and look-and-feel issues
from dialogue and abstract presentation decisions via mechanisms employing a client-
server architecture. To separate application code from user interf ...

**Keywords**: artists, client-server, concurrency, event-based integration, user interface
architectures

**12** Manchester data-flow: a progress report
J. R. Gurd, D. F. Snelling
August 1992 **Proceedings of the 6th international conference on Supercomputing ICS '92**
Publisher: ACM Press

Full text available: 📄 pdf(1.23 MB)       Additional Information: full citation, abstract, references, citings, index terms

The Manchester Data-Flow Machine, MDFM, has evolved continuously during the past decade. By the time the prototype uniprocessor hardware system was decommissioned, in 1989, the putative multi-processor architecture comprised separate Processing Elements and Structure Store Units, together with a "global" Allocator and Throttle Unit, all linked by a packet-based Interconnection Switch. The decisions leading to this design are well documented in the literature, but some ...

**13** Balancing power consumption in multiprocessor systems
Andreas Merkel, Frank Bellosa
April 2006 **ACM SIGOPS Operating Systems Review , Proceedings of the 2006 EuroSys conference EuroSys '06**, Volume 40 Issue 4
Publisher: ACM Press

Full text available: 📄 pdf(688.06 KB)    Additional Information: full citation, abstract, references, index terms

Actions usually taken to prevent processors from overheating, such as decreasing the frequency or stopping the execution flow, also degrade performance. Multiprocessor systems, however, offer the possibility of moving the task that caused a CPU to overheat away to some other, cooler CPU, so throttling becomes only a last resort taken if all of a system's processors are hot. Additionally, the scheduler can take advantage of the energy characteristics of individual tasks, and distribute hot tasks ...

**Keywords**: energy estimation, energy-aware scheduling, event counters, task migration, thermal management

**14** Combining optimism limiting schemes in time warp based parallel simulations
Kevin Jones, Samir R. Das
December 1998 **Proceedings of the 30th conference on Winter simulation WSC '98**
Publisher: IEEE Computer Society Press

Full text available: 📄 pdf(70.94 KB)    Additional Information: full citation, references, citings, index terms

**15** Disk Drive Roadmap from the Thermal Perspective: A Case for Dynamic Thermal Management
Sudhanva Gurumurthi, Anand Sivasubramaniam, Vivek K. Natarajan
May 2005 **ACM SIGARCH Computer Architecture News , Proceedings of the 32nd Annual International Symposium on Computer Architecture ISCA '05**, Volume 33 Issue 2
Publisher: IEEE Computer Society, ACM Press

Full text available: 📄 pdf(243.57 KB)    Additional Information: full citation, abstract, index terms

The importance of pushing the performance envelope of disk drives continues to grow, not just in the server market but also in numerous consumer electronics products. One of the most fundamental factors impacting disk drive design is the heat dissipation and its effect on drive reliability, since high temperatures can cause off-track errors, or even head crashes. Until now, drive manufacturers have continued to meet the 40% annual growth target of the internal data rates (IDR) by increasing RPMs ...

**16** Thermal issues in disk drive design: Challenges and possible solutions
Sudhanva Gurumurthi, Anand Sivasubramaniam
February 2006

**ACM Transactions on Storage (TOS)**, Volume 2 Issue 1
**Publisher:** ACM Press
Full text available: pdf(2.14 MB)     Additional Information: full citation, abstract, references, index terms

The importance of pushing the performance envelope of disk drives continues to grow in the enterprise storage market. One of the most fundamental factors impacting disk drive design is heat dissipation, since it directly affects drive reliability. Until now, drive manufacturers have continued to meet the 40&percnt; annual growth target of the internal data-rates (IDR) by increasing RPMs and shrinking platter sizes, both of which have counteracting effects on the heat dissipation within a drive. ...

**Keywords**: Disk drive, technology scaling, thermal management

**17** Tolerating memory latency through push prefetching for pointer-intensive applications ☐

Chia-Lin Yang, Alvin R. Lebeck, Hung-Wei Tseng, Chien-Hao Lee
December 2004 **ACM Transactions on Architecture and Code Optimization (TACO)**,
Volume 1 Issue 4
**Publisher:** ACM Press
Full text available: pdf(590.24 KB)     Additional Information: full citation, abstract, references, index terms

Prefetching is often used to overlap memory latency with computation for array-based applications. However, prefetching for pointer-intensive applications remains a challenge because of the irregular memory access pattern and pointer-chasing problem. In this paper, we proposed a cooperative hardware/software prefetching framework, the push architecture, which is designed specifically for linked data structures. The push architecture exploits program structure for future address generation instea ...

**Keywords**: Prefetch, linked data structures, memory hierarchy, pointer-chasing

**18** Scheduling computations on a software-based router ☐

Xiaohu Qie, Andy Bavier, Larry Peterson, Scott Karlin
June 2001 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2001 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '01**, Volume 29 Issue 1
**Publisher:** ACM Press
Full text available: pdf(1.46 MB)     Additional Information: full citation, abstract, references, citings

Recent efforts to add new services to the Internet have increased the interest in software-based routers that are easy to extend and evolve. This paper describes our experiences implementing a software-based router, with a particular focus on the main difficulty we encountered: how to schedule the router's CPU cycles. The scheduling decision is complicated by the desire to differentiate the level of service for different packet flows, which leads to two fundamental conflicts: (1) assigning proce ...

**19** Ensemble-level Power Management for Dense Blade Servers ☐

Parthasarathy Ranganathan, Phil Leech, David Irwin, Jeffrey Chase
May 2006 **ACM SIGARCH Computer Architecture News , Proceedings of the 33rd annual international symposium on Computer Architecture ISCA '06**, Volume 34 Issue 2
**Publisher:** IEEE Computer Society, ACM Press
Full text available: pdf(333.39 KB)     Additional Information: full citation, abstract, index terms

One of the key challenges for high-density servers (e.g., blades) is the increased costs in addressing the power and heat density associated with compaction. Prior approaches have mainly focused on reducing the heat generated at the level of an individual server. In contrast, this work proposes power efficiencies at a larger scale by leveraging statistical properties of concurrent resource usage across a collection of systems ("ensemble"). Specifically, we discuss an implementation of this appro ...

**20**

### Beyond objects: a software design paradigm based on process control

Mary Shaw

January 1995 **ACM SIGSOFT Software Engineering Notes**, Volume 20 Issue 1

**Publisher:** ACM Press

Full text available: pdf(1.05 MB)    Additional Information: full citation, abstract, citings, index terms

A standard demonstration problem in object-oriented programming is the design of an automobile cruise control. This design exercise demonstrates object-oriented techniques well, but it does not ask whether the object-oriented paradigm is the best one for the task. Here we examine the alternative view that cruise control is essentially a control problem. We present a new software organization paradigm motivated by process control loops. The control view leads us to an architecture that is dominat ...

Results 1 - 20 of 200          Result page: **1**  2  3  4  5  6  7  8  9  10    next